

Building Transparent User Gateways at NERSC

Shreyas Cholia
NERSC/LBL

GlobusWorld, March 2nd 2010
Argonne National Laboratory



What is NERSC

- **National Energy Research Scientific Computing Center**
 - DOE Office of Science User Facility at Lawrence Berkeley National Laboratory
- **Mission:**
 - accelerate the pace of scientific discovery in the DOE Office of Science community by providing high-performance computing, information, data, and communications services.



NERSC is the Primary Computing Facility for the Office of Science

- NERSC serves a large population

Approximately 3000 users,
400 projects, 500 code instances


- Focus on “unique” resources

- High end computing systems
- High end storage systems
 - File system and tape archive
- Interface to high speed networking

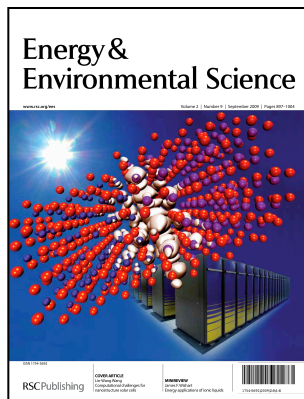
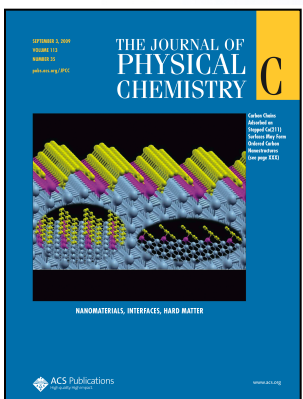
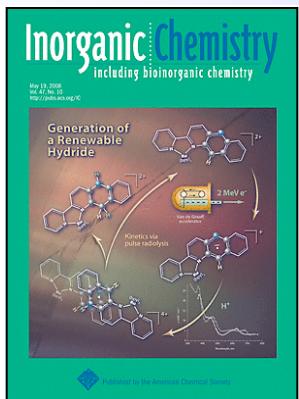

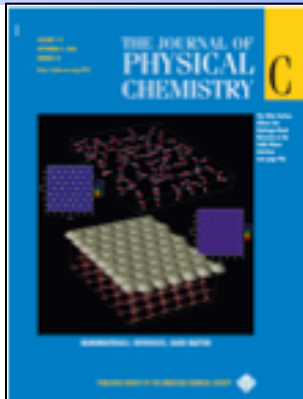

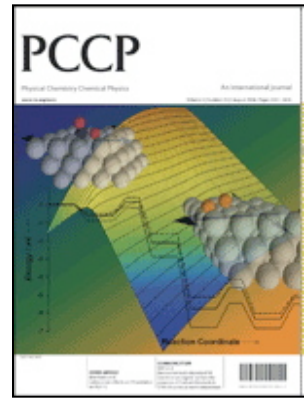
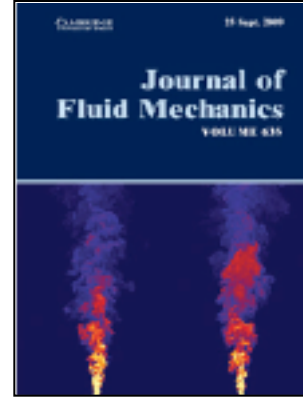



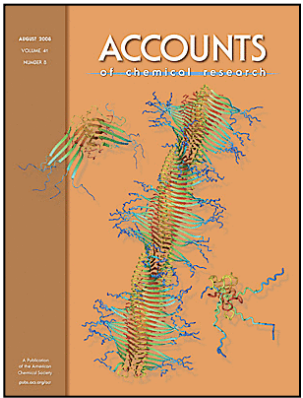
- Science-driven

- Science problems used in machine procurements and performance metrics
- Science services

2009 Allocations



High quality science results from simulations of many different scales

 <p>Energy & Environmental Science</p>	 <p>THE JOURNAL OF PHYSICAL CHEMISTRY C</p>	 <p>Inorganic Chemistry</p>	 <p>JCTC</p>	 <p>THE JOURNAL OF PHYSICAL CHEMISTRY C</p>	 <p>THE JOURNAL OF PHYSICAL CHEMISTRY B</p>
64-6000 cores	4-128 cores	32 cores	1100-4600 cores	32-1024 cores	256 cores
 <p>PCCP</p>	 <p>Journal of Fluid Mechanics</p>	 <p>PHYSICS TODAY</p>	 <p>THE JOURNAL OF PHYSICAL CHEMISTRY B</p>	 <p>nature</p>	 <p>ACCOUNTS of chemical research</p>
32-256 cores	64-4096 cores	1000-4000 cores	80-260 cores	64-4700 cores	1100-4600 cores

NERSC Systems

PDSF
HEP
Linux
Cluster

Franklin
Cray XT4

Hopper
Cray XT5

DaVinci
SGI Altix

Carver
IBM
IDataplex
Cluster

Magellan
IBM
Idataplex
Cloud

GPFS Global FS

**HPSS Mass
Storage**



Diversity of Users and Systems

- **Users have differing application requirements**
- **Wide range of access patterns**
- **Multiple systems to meet different user needs**

Grid Usage

- **Open Science Grid has an allocation at NERSC – primarily for MPI users**
- **PDSF – STAR, ATLAS, ALICE, Planck using Grid tools to access system**
- **GridFTP – very popular for WAN data movement**
- **Grid certificate based workflow – very unpopular among wider user community ☹**
 - Managing certificates and dealing with existing CLI tools a huge deterrent for new users

GSI is still Useful

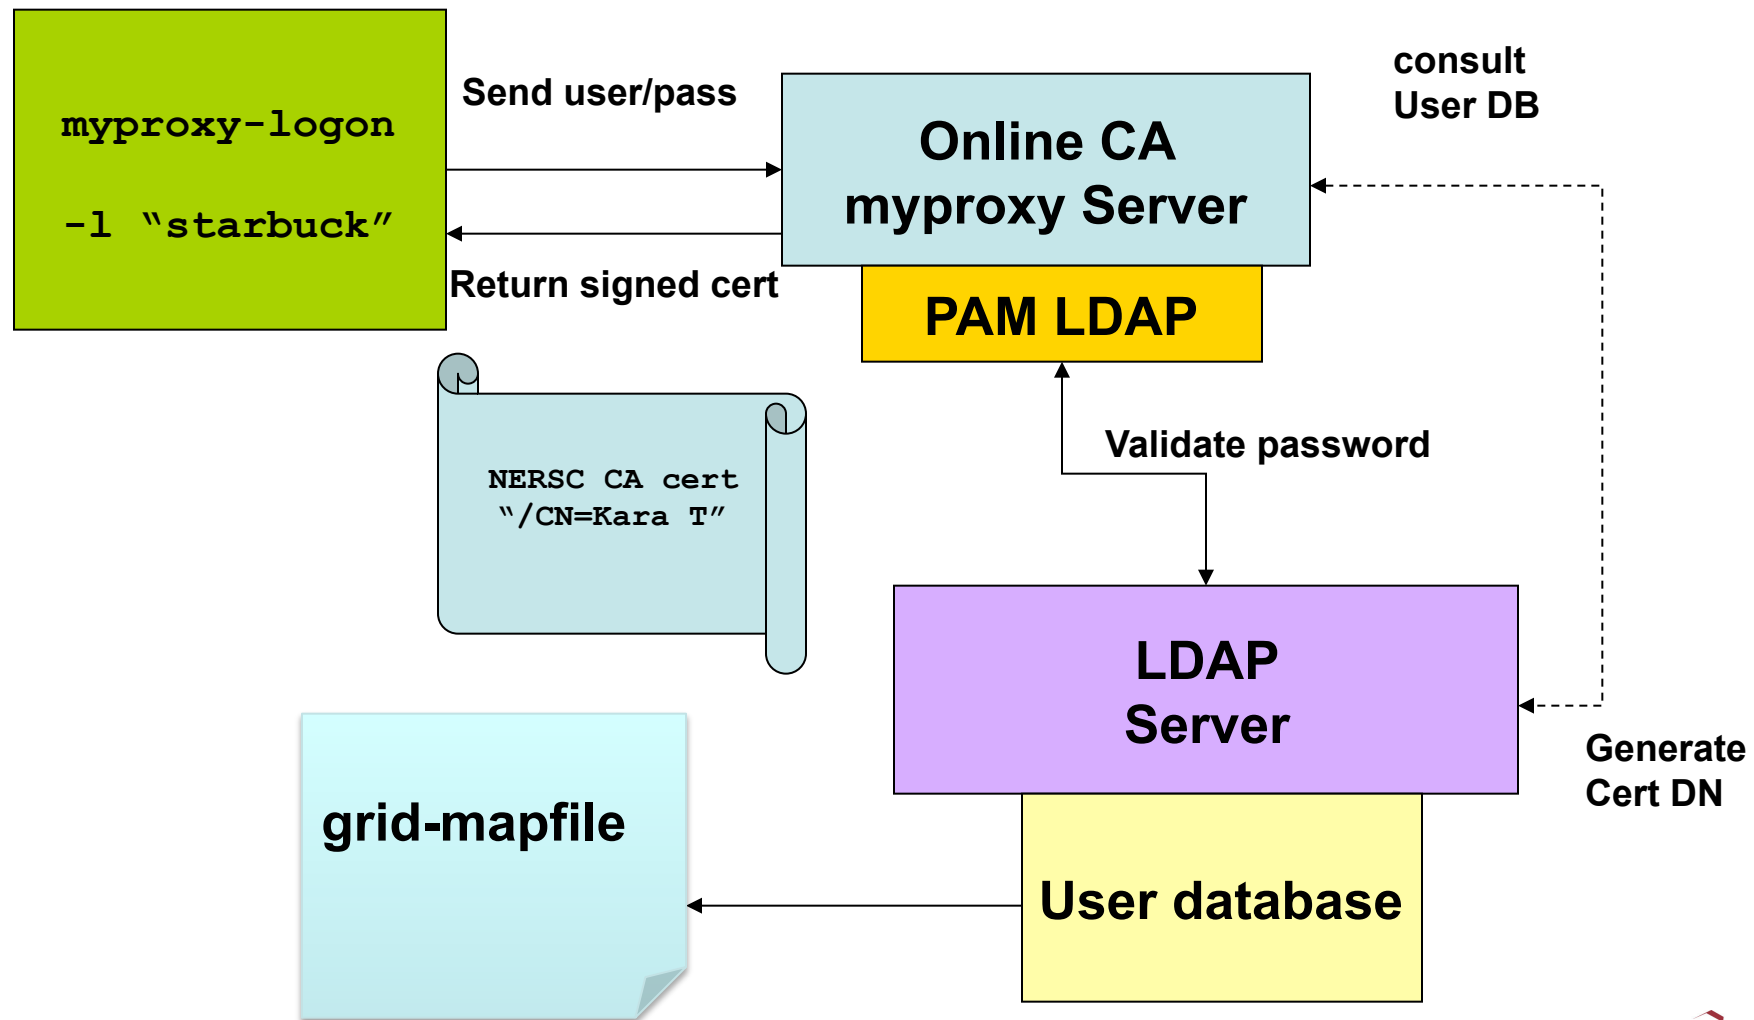
- **As site admins and developers we see the value of the GSI infrastructure**
 - Provides a secure SSO to HPC resources
 - Common credential across multiple resources
 - Reduces password based access
 - Separates AuthN from AuthZ
 - Preserves user identity
- **How do we preserve these semantics without burdening the user?**



Step 1: Simplify Certificate Generation and Management

- **NERSC runs MyProxy in CA mode**
 - Single command to get a cert
 - Generates a brand new Short-lived cert based on username/password authentication
 - Backended to existing NERSC LDAP database
- **Generate persistent DNs AND populate grid-mapfiles on various systems using the same LDAP info**
 - Include example

NERSC CA Service

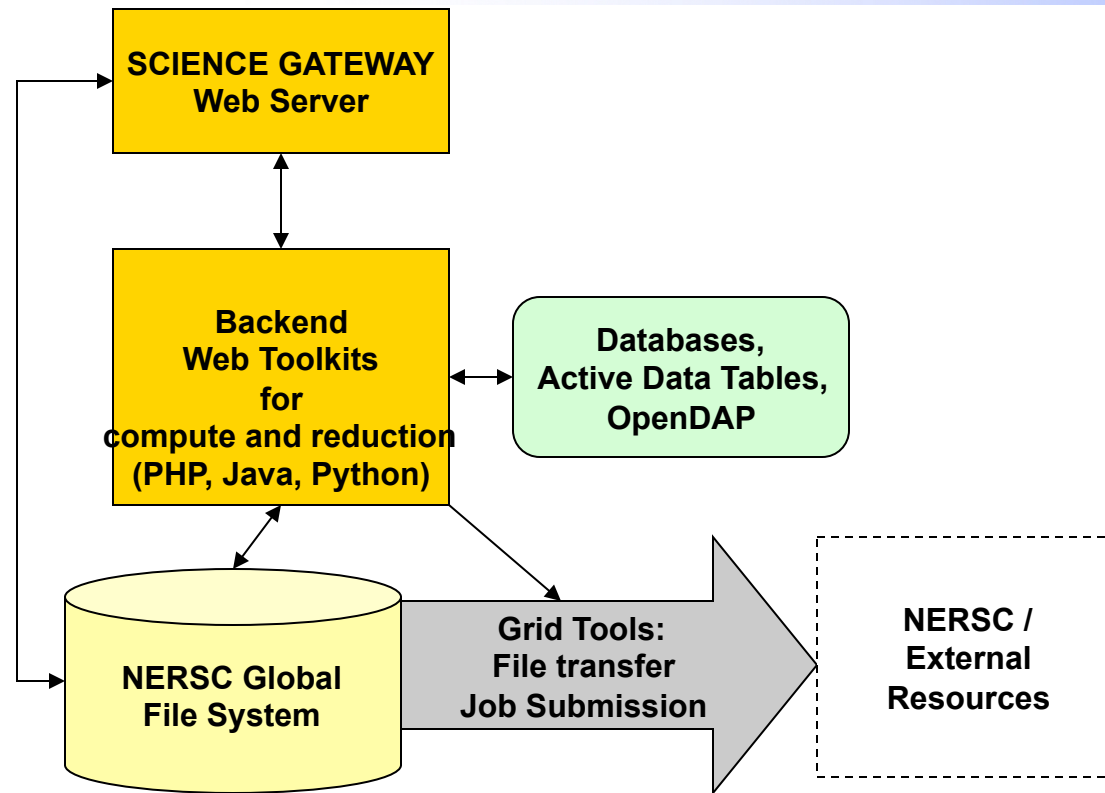




Step 2: Hide Complexity through Web Gateways

- **Web access methods to NERSC resources**
- **Users very comfortable with web paradigm. Now expect it for usability**
- **HPC should be as easy online-banking**
- **Custom gateways for individual science projects**
 - Users don't want to see generic options or tools not applicable to their science
 - Users don't want to deal with grid centric views
- **NERSC gateway services**
 - host the gateway
 - assist in building the webapp
 - provide building blocks to science groups for their own apps.

NERSC Science Gateways



Provide sophisticated building blocks for science on the web

- Direct access to massive data (NGF) through web-server
- Reduction and compute backend tools available (PHP, Python, Java)
- Databases, active tables, OpenDAP, File Browsing tools
- Grid tools for file transfer and job submission

Planck Gateway - Login

Computational Cosmology Center


BERKELEY LAB COMPUTING SCIENCES



Lawrence Berkeley National Laboratory operates this computer system under contract to the U.S. Department of Energy. This computer system is the property of the United States Government and is for authorized use only. Users (authorized or unauthorized) have no explicit or implicit expectation of privacy.

Any or all uses of this system and all files on this system may be intercepted, monitored, recorded, copied, audited, inspected, and disclosed to site, Department of Energy, and law enforcement personnel, as well as authorized officials of other agencies, both domestic and foreign. By using this system, the user consents to such interception, monitoring, recording, copying, auditing, inspection, and disclosure at the discretion of authorized site or Department of Energy personnel.

Unauthorized or improper use of this system may result in administrative disciplinary action and civil and criminal penalties. By continuing to use this system you indicate your awareness of and consent to these terms and conditions of use. LOG OFF IMMEDIATELY if you do not agree to the conditions stated in this warning.



Username:

Password:

Login




U.S. DEPARTMENT OF
ENERGY

Office of
Science



Planck Gateway – Web Application



Computational Cosmology Center

BERKELEY LAB COMPUTING SCIENCES

Multimod is an executable within the Level-S suite of Planck simulation tools. Its purpose is to produce timestreams from various input signals (sky maps, point sources, 1/f noise, sorption cooler noise, etc).

This page allows the user to create task-list files and job setup scripts which may be used to run simulations using mpiBatch.

Please enter required parameters and select a set of detectors.

Required Parameters

Simulation Name:	Timestream Name:
<input style="width: 90%;" type="text"/>	<input style="width: 90%;" type="text"/>

Machine	Number of Cores:
<div style="border: 1px solid black; padding: 2px;"> Franklin </div>	<input style="width: 90%;" type="text"/>

Focal Plane Database:

-- Choose a File --

Satellite Pointing File:

-- Choose a File --

Benefits

- **Perform actions on as the user on remote systems.**
- **User doesn't ever know about backend globus/GSI layer**
- **We preserve authorization at the infrastructure layer by using GSI certificates to access resources**

Currently

- **Set of ad-hoc portals**
 - DeepSky
 - Planck
 - GCRM (demo)
 - QCD
 - OpenDAP service for NETCDF data



Step 3: Access Resources through Reusable Web API

- Ad hoc approach has been productive but we see common patterns
- Encapsulate these patterns as building blocks for Science Gateways
- Building block API should be very easy to invoke eg. via a simple web page
 - Every resource should be encapsulated as a URL with a simple set of associated actions

REST

- **Representational State Transfer**
- **Every resource is represented by a unique http URL**
- **Actions are defined by standard HTTP methods: GET, POST, PUT, DELETE**
- **Think of it as an API that uses the language of HTTP**
- **NERSC Web Toolkit (NEWT) - RESTful service that provides access to NERSC resources**
- **NEWT combines NERSC database resources, Grid resources and other RESTful services under a single API**

NEWT API examples

VERB	RESOURCE	DESCRIPTION
POST	/resource/job/	submit POST data to queue on R, return job id
GET	/resource/file/path/fname	get "fname" in "path" on R, copy it to apache server and download the file
GET	/user/username	get user account info

We are building this using the Django python framework. Authentication (MyProxy), File Transfer (GridFTP) and Job Management (GRAM) is completely transparent to the user and web application developer

HPC as a Service

- **Build web applications through REST API**
- **No need for science team to learn underlying framework**
- **GSI and globus calls essentially act at the backend protocol layer**
- **User interacts with a web application that exposes the necessary components of the underlying application**

In Development

- **Build reusable components and expose them through a REST API.**
- **Globus.org is very interesting this context**
 - Will provide us with a backend service layer to perform distributed computing tasks
 - Waiting for RESTful interface to service



Thanks

Josh Boverhof

David Skinner

R.K. Owen

Contact: Shreyas Cholia

scholia@lbl.gov

Questions?

